Vehicles

Experiments in Synthetic Psychology

Valentino Braitenberg

The MIT Press
Cambridge, Massachusetts
London, England

Introduction

Let the Problem of the Mind Dissolve in Your Mind

This is an exercise in fictional science, or science fiction, if you like that better. Not for amusement: science fiction in the service of science. Or just science, if you agree that fiction is part of it, always was, and always will be as long as our brains are only minuscule fragments of the universe, much too small to hold all the facts of the world but not too idle to speculate about them.

I have been dealing for many years with certain structures within animal brains that seemed to be interpretable as pieces of computing machinery because of their simplicity and/or regularity. Much of this work is only interesting if you are yourself involved in it. At times, though, in the back of my mind, while I was counting fibers in the visual ganglia of the fly or synapses in the cerebral cortex of the mouse, I felt knots untie, distinctions dissolve, difficulties disappear, difficulties I had experienced much earlier when I still held my first naive philosophical approach to the problem of the mind. This process of purification has been, over the years, a delightful experience. The text I want you to read is designed to convey some of this

to you, if you are prepared to follow me not through a world of real brains but through a toy world that we will create together.

We will talk only about machines with very simple internal structures, too simple in fact to be interesting from the point of view of mechanical or electrical engineering. Interest arises, rather, when we look at these machines or vehicles as if they were animals in a natural environment. We will be tempted, then, to use psychological language in describing their behavior. And yet we know very well that there is nothing in these vehicles that we have not put in ourselves. This will be an interesting educational game.

Our vehicles may move in water by jet propulsion. Or you may prefer to imagine them moving somewhere between galaxies, with negligible gravitational pull. Remember, however, that their jets must expel matter in order to function at all, and this implies replenishment of the food stores within the vehicles, which might be a problem between galaxies. This suggests vehicles moving on the surface of the earth through an agricultural landscape where they have good support and can easily find the food or fuel they need. (Indeed the first few chapters here conjure up images of vehicles swimming around in the water, while later what comes to mind are little carts moving on hard surfaces. This is no accident, if the evolution of vehicles 1 to 14 in any way reflects the evolution of animal species.)

It does not matter. Get used to a way of thinking in which the hardware of the realization of an idea is much less important than the idea itself. Norbert Wiener was emphatic about this when he formulated the title of his famous book: *Cybernetics, or Control and Communication in Animals and Machines.*

Vehicle 1

Getting Around

Vehicle 1 is equipped with one sensor and one motor (figure 1). The connection is a very simple one. The more there is of the quality to which the sensor is tuned, the faster the motor goes. Let the quality be temperature and let the force exerted by the motor be exactly proportionate to the absolute temperature (the temperature above zero degrees Kelvin) measured by the sensor. The vehicle will move, wherever it is (the absolute temperature is nowhere equal to zero), in the direction in which it happens to be pointing. It will slow down in cold regions and speed up where it is warm.

Here we have introduced a bit of Aristotelian physics. Aristotle, like everybody else between this ancient Greek philosopher and the less ancient Italian physicist Galileo, thought that the speed of a moving body is proportionate to the force that drives it. This is true in most instances, namely when there is friction to slow down the vehicle. Normally friction will see to it that the velocity becomes zero in the absence of any force, that it will stay at a certain small value for a certain small force, at a higher value for a higher force, and so forth.

Of course, as you all know, this is not true for heavenly bodies

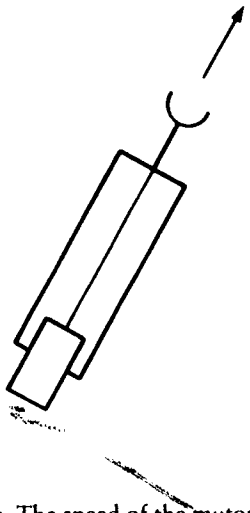

Figure 1
Vehicle 1, the simplest vehicle. The speed of the motor (rectangular box at the tail end) is controlled by a sensor (half circle on a stalk, at the front end). Motion is always forward, in the direction of the arrow, except for perturbations.

(especially if you don't invest astronomical time in observing them). Their velocity is a complicated result of all the forces that ever hit them. This is another reason for letting our vehicles move in water or on the surface of the earth rather than in outer space.

In this Aristotelian world our vehicle number 1 may even come to rest. This will happen when it enters a cold region where the force exerted by its motor, being proportionate to the temperature, becomes smaller than the frictional force.

Once you let friction come into the picture, other amazing things may happen. In outer space Vehicle 1 would move on a straight course with varying speed (the gravitational pull of neighboring galaxies averages out to nothing). Not so on earth. The friction, which is nothing but the sum of all the microscopic forces that arise in a situation too messy to be analyzed in detail, may not be quite symmetrical. As the vehicle pushes forward against frictional forces, it will deviate from its course. In the long run it will be seen to move in a complicated trajectory, curving one way or the other without apparent good reason. If it is very small, its motion will be quite erratic, similar to "Brownian motion," only with a certain drive added.

Imagine, now, what you would think if you saw such a vehicle swimming around in a pond. It is restless, you would say, and does not like warm water. But it is quite stupid, since it is not able to turn back to the nice cold spot it overshot in its restlessness. Anyway, you would say, it is ALIVE, since you have never seen a particle of dead matter move around quite like that.

Vehicle 2

Fear and Aggression

Vehicle 2 is generally similar to Vehicle 1 except that it has two sensors, one on each side, and two motors, right and left (figure 2). You may think of it as being a descendant of Vehicle 1 through some incomplete process of biological reduplication: two of the earlier brand stuck together side by side. Again, the more the sensors are excited, the faster the motors run.

Of course you notice right away that we can make three kinds of such vehicles, depending on whether we connect (a) each sensor to the motor on the same side, (b) each sensor to the motor on the opposite side, or (c) both sensors to both motors. We can immediately dismiss case (c), for this is nothing but a somewhat more luxurious version of Vehicle 1. The difference between (a) and (b), however, is very interesting.

Consider (a) first. This vehicle will spend more time in the places where there is less of the stuff that excites its sensors and will speed up when it is exposed to higher concentrations. If the source of the stuff (say, light in the case of light sensors) is directly ahead, the vehicle may hit the source unless it is deflected from its course. If the source is to one side (figure 3), one of the sensors, the one nearer to the source, is excited more than the other. The corresponding

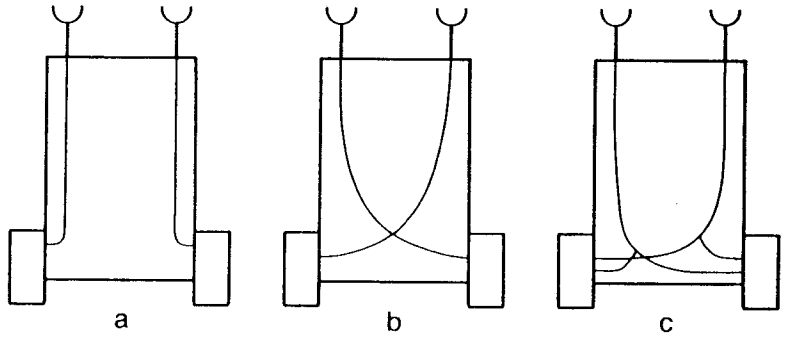

Figure 2
Vehicle 2, with two motors and two sensors; otherwise like Vehicle 1. The connections differ in a, b, and c.

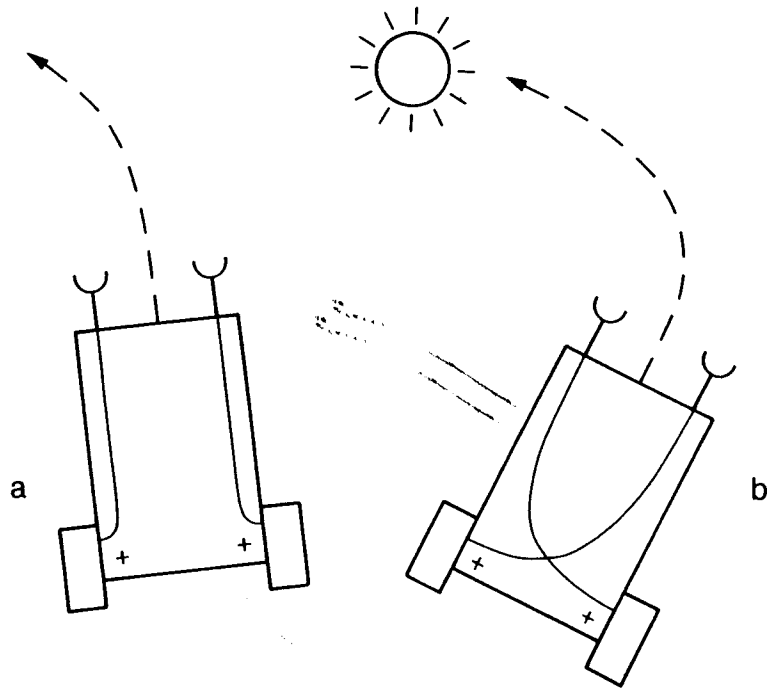

Figure 3
Vehicles 2a and 2b in the vicinity of a source (circle with rays emanating from it). Vehicle 2b orients toward the source, 2a away from it.

motor will work harder. And as a consequence the vehicle will turn away from the source.

Now let us try the other scheme of sensory-motor connections, (b) in figure 3. No change if the source is straight ahead. If it is to one side, however, we notice a difference with respect to Vehicle 2a. Vehicle 2b will turn toward the source and eventually hit it. There is no escaping: as long as 2b stays in the vicinity of the source, no matter how it stumbles and hesitates, it will hit the source frontally in the end. Only in the unlikely case that a strong perturbation in its course makes it turn exactly away from the source, and no further perturbation occurs, can it escape its fate.

Let Vehicles 2a and 2b move around in their world for a while and watch them. Their characters are quite opposite. Both DISLIKE sources. But 2a becomes restless in their vicinity and tends to avoid them, escaping until it safely reaches a place where the influence of the source is scarcely felt. Vehicle 2a is a COWARD, you would say. Not so Vehicle 2b. It, too, is excited by the presence of sources, but resolutely turns toward them and hits them with high velocity, as if it wanted to destroy them. Vehicle 2b is AGGRESSIVE, obviously.

Vehicle 3

Love

The violence of Vehicle 2b, no less than the cowardice of its companion 2a, are traits that call for improvement. There is something very crude about a vehicle that can only be excited by the things it smells (or sees or feels or hears) and knows no soothing or relaxing stimuli. What comes to mind is to introduce some inhibition in the connections between the sensors and the motors, switching the sign of the influence from positive to negative. This will let the motor slow down when the corresponding sensor is activated. Again we can make two variants, one with straight and one with crossed connections (figure 4). Both will slow down in the presence of a strong stimulus and race where the stimulus is weak. They will therefore spend more time in the vicinity of the source than away from it. They will actually come to rest in the immediate vicinity of the source.

But here we notice a difference between the vehicle with straight connections and the one with crossed connections. Approaching the source, the first (figure 4a) will orient toward it, since on an oblique course the sensor nearer to the source will slow down the motor on the same side, producing a turn toward that side. The vehicle with straight connections will come to rest facing the

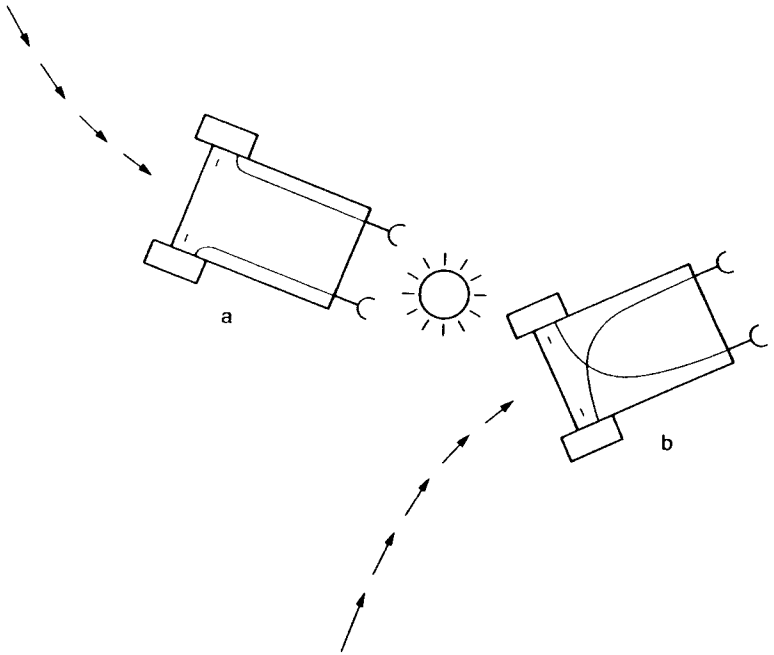

Figure 4
Vehicle 3, with inhibitory influence of the sensors on the motors.

source. The vehicle with crossed connections (figure 4b) for analogous reasons will come to rest facing away from the source and may not stay there very long, since a slight perturbation could cause it to drift away from the source. This would lessen the source's inhibitory influence, causing the vehicle to speed up more and more as it gets away.

You will have no difficulty giving names to this sort of behavior. These vehicles LIKE the source, you will say, but in different ways. Vehicle 3a LOVES it in a permanent way, staying close by in quiet admiration from the time it spots the source to all future time. Vehicle 3b, on the other hand, is an EXPLORER. It likes the nearby source all right, but keeps an eye open for other, perhaps stronger sources, which it will sail to, given a chance, in order to find a more permanent and gratifying appeasement.

But this is not yet the full development of Vehicle 3. We are now ready to make a more complete model using all the behavioral traits at our disposal. Call it Vehicle 3c. We give it not just one pair of sensors but four pairs, tuned to different qualities of the environment, say light, temperature, oxygen concentration, and amount of organic matter (figure 5). Now we connect the first pair to the motors with uncrossed excitatory connections, as in Vehicle 2a, the second pair with crossed excitatory connections, as in Vehicle 2b, and the third and the fourth pairs with inhibitory connections, crossed and uncrossed, as in Vehicles 3b and 3a.

This is now a vehicle with really interesting behavior. It dislikes high temperature, turns away from hot places, and at the same time seems to dislike light bulbs with even greater passion, since it turns toward them and destroys them. On the other hand it definitely seems to prefer a well-oxygenated environment and one containing many organic molecules, since it spends much of its time in such places. But it is in the habit of moving elsewhere when the supply of either organic matter or (especially) oxygen is low. You cannot help admitting that Vehicle 3c has a system of VALUES, and, come to

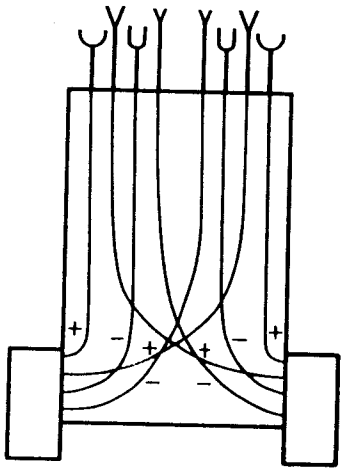

Figure 5
A multisensorial vehicle of brand 3c.

think of it, KNOWLEDGE, since some of the habits it has, like destroying light bulbs, may look quite knowledgeable, as if the vehicle knows that light bulbs tend to heat up the environment and consequently make it uncomfortable to live in. It also looks as if it knows about the possibility of making energy out of oxygen and organic matter because it prefers places where these two commodities are available.

But, you will say, this is ridiculous: knowledge implies a flow of information from the environment into a living being or at least into something like a living being. There was no such transmission of information here. We were just playing with sensors, motors, and connections: the properties that happened to emerge may look like knowledge but really are not. We should be careful with such words.

You are right. We will explain in a later chapter (on Vehicle 6) how knowledge may enter a system of connections. And we will introduce an alternative way of incorporating knowledge into the system in our chapter on Vehicle 7. In any case, once knowledge is incorporated, the resulting vehicle may look and behave quite like our Vehicle 3c.

Meanwhile I invite you to consider the enormous wealth of different properties that we may give Vehicle 3c by choosing various sensors and various combinations of crossed and uncrossed, excitatory and inhibitory, connections.

If you consider the possibility of strong and weak influences from the sensors to the motors, you realize that the variety becomes even greater. The vehicle may not care much about light but care very much about temperature. Its sense of smell may be much keener for organic matter than it is for oxygen or vice versa. And there may be many more than just four pairs of sensors and four sensory qualities: the vehicles may be equipped with all sorts of shrewd detectors of energy and of chemicals. But this is best discussed in connection with a new idea incorporated in the vehicles of the next chapter.

Vehicle 4

Values and Special Tastes

We are now in a position to create a new brand of vehicle, starting from all the varieties of Vehicle 3, by working on the connections between sensors and motors. They were, up to now, of two very simple kinds: the more the sensor was excited, the faster the corresponding motor ran, or, alternatively, the more the sensor was excited, the slower the motor ran. We did not care what the rules of the dependence were, as long as they were of the nature "the more, the more" or "the more, the less." The vast class of mathematical functions describing such dependences is sometimes called monotonic. Obviously, there is something very simple-minded about creatures governed by such unconditioned likes or dislikes, and we can easily see how such the-more-the-merrier behavior could lead to disaster. Think what happens in the case of a tendency to follow downhill slopes!

Let us consider the following improvement. The activation of a certain sensor will make the corresponding motor run faster, but only up to a point, where the speed of the motor reaches a maximum. Beyond this point, if the sensor is activated even more strongly, the speed will decrease again (figure 6). The same sort of dependence, with a maximum efficiency at a certain level of sensor

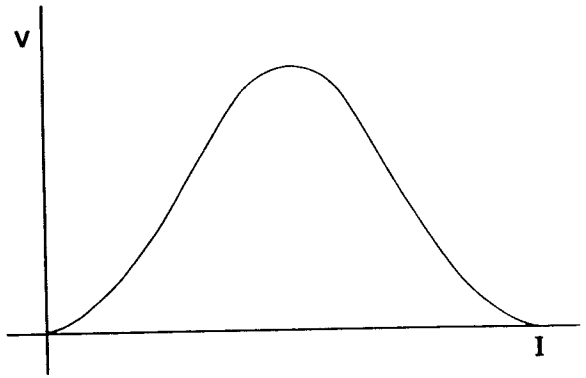

Figure 6
A nonlinear dependence of the speed of the motor V on the intensity of stimulation I, with a maximum for a certain intensity.

activation, can be engineered for the inhibitory connections between sensor and motor. We may set the maximum efficiency of the various sensors at any level we choose, and we may even play with dependences having more than one maximum. Any vehicle constructed according to this prescription we will assign to a new brand, labeled 4a. Of course, if you like, you can keep some of the connections of the old monotonic type and mix them with the nonmonotonic ones in every possible combination.

You will have a hard time imagining the variety of behavior displayed by the vehicles of brand 4a. A 4a vehicle might navigate toward a source (as Vehicle 2b would) and then turn away when the stimulus becomes strong, circle back and then turn away over and over again, perhaps describing a trajectory in the form of a figure eight. Or it might orbit around the source at a fixed distance, like a satellite around the earth, its course being corrected toward the source by a weaker stimulus and away from the source by a stronger stimulus, depending on whether the stimulus intensity is

on one side or the other of the maximum describing the sensory-motor dependence (figure 7). Vehicle 4a might like one sort of stimulus when it is weak but not when it is too strong; it might like another stimulus better the stronger it becomes. It might turn away from a weak smell and destroy the source of a strong one. It might visit in alternation a source of smell and a source of sound, turning away from both with a change of temperature.

Watching vehicles of brand 4a in a landscape of sources, you will be delighted by their complicated trajectories. And I am sure you will feel that their motives and tastes are much too varied and intricate to be understood by the observer. These vehicles, you will say, are governed by INSTINCTS of various sorts and, alas, we just don't know how Nature manages to embody instincts into a piece of brain.

You forget, of course, that we have ourselves designed these vehicles.

But instincts are a lowly sort of behavior anyway. We can do better. Let us improve on type 4a by adding a new sort of connection between sensors and motors. This time the influence of the sensor on the motor is no longer smooth; there are definite breaks. There might be a range of intensities of sensory stimulation for which the motor is not activated at all and then, under stronger stimuli, the motors are running at full speed. Or else, there might be smooth changes of motor activation for certain ranges, with abrupt changes in between. A very lifelike pattern would be: no activation up to a threshold value of the stimulus, and increasing activation beyond the threshold, starting with a certain fixed minimum (figure 8). You are by now experienced in the art of creative invention and will have no difficulty dreaming up more schemes of this sort.

In a way these new vehicles, which we call 4b, are already contained in the vast class of vehicles 4a, since abruptness of behavior can of course be simulated with any degree of approximation by functional dependences that are in reality, mathematically speak-

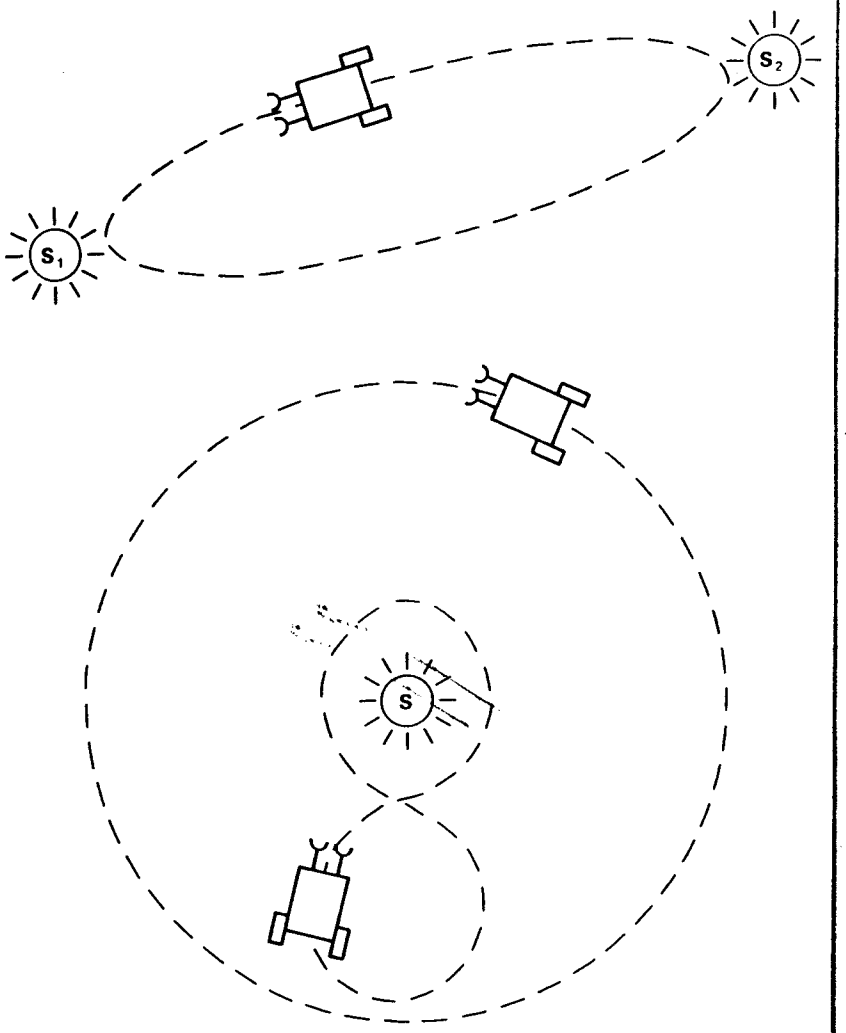

Figure 7
Trajectories of vehicles of brand 4a around or between sources.

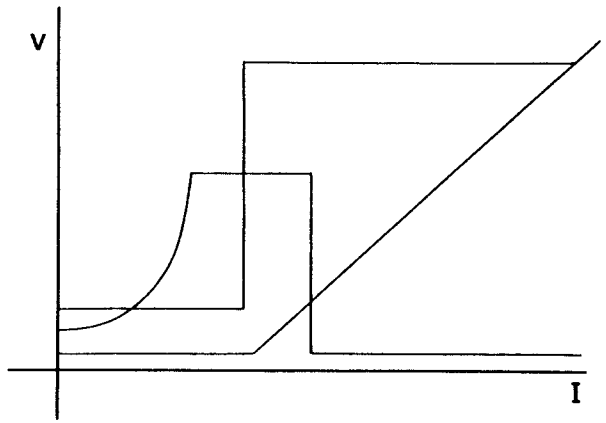

Figure 8
Various bizarre kinds of dependence of the speed of the motor (ordinate) on the intensity of stimulation (abscissa) in Vehicle 4b.

ing, continuous. Moreover, if friction plays a role, as we have already decided it should, thresholds in motor activation would ensue naturally: the vehicle will start moving only when the force exerted by the motor exceeds a certain value, sufficient to overcome the initial friction.

Whatever their origin, thresholds in some behavior patterns make a lot of difference in the eye of the observer. These creatures, the observer would say, ponder over their DECISIONS. When you come close to them with a lure, it takes them some time to get going. Yet once they have decided, they can act quite quickly. They do indeed seem to act in a spontaneous way: none of this passive being attracted one way or the other that was so obvious in the vehicles of the more lowly types. You would almost be tempted to say: where decisions are being made, there must be a WILL to make them. Why not? For all we know, this is not the worst criterion for establishing the existence of free will.

Vehicle 5

Logic

At this point we are ready to make a fundamental discovery. We have gathered evidence for what I would like to call the "law of uphill analysis and downhill invention." What I mean is this. It is pleasurable and easy to create little machines that do certain tricks. It is also quite easy to observe the full repertoire of behavior of these machines—even if it goes beyond what we had originally planned, as it often does. But it is much more difficult to start from the outside and to try to guess internal structure just from the observation of behavior. It is actually impossible in theory to determine exactly what the hidden mechanism is without opening the box, since there are always many different mechanisms with identical behavior. Quite apart from this, analysis is more difficult than invention in the sense in which, generally, induction takes more time to perform than deduction: in induction one has to search for the way, whereas in deduction one follows a straightforward path.

A psychological consequence of this is the following: when we analyze a mechanism, we tend to overestimate its complexity. In the uphill process of analysis, a given degree of complexity offers more resistance to the workings of our mind than it would if we encountered it downhill, in the process of invention. We have already seen this happen when the observer of Vehicle 4b conjectured that the vehicle does some thinking before it reaches a decision, suggesting complicated internal processes where in reality there was nothing but a threshold device waiting for sufficient activation. The patterns of behavior described in the vehicles of type 4a undoubtedly suggest much more complicated machinery than that which was actually used in designing them.

We may now take pleasure in this and create simple "brains" for our vehicles, which will indeed (as experience shows) tax the mind of even the most playful analyst. All we have to do is introduce special elements, called threshold devices, which will be either interposed between sensors and motors or connected to each other in complexes that receive some input from the sensors and give some output to the motors.

The individual threshold device is of the simplest sort: it gives no output if its input line carries a signal below the threshold, and it gives full output beyond the threshold. We will also use another variety giving output all the time unless the input carries a signal above the threshold. Each of these devices is fitted with a knob which may be turned to set the threshold, so that the input would become effective with one, two, or any specified number of input activation units. (The word threshold of course implies that, for a given threshold value, any input stronger than the one specified would also be effective.)

We are not limited to the types of connections through which the threshold devices activate each other. We can also use another kind, call them "inhibitory," which counteract the activation that comes from other sources (figure 9).

In order to make a brain out of threshold devices, we may connect them together one to one, or many to one, or one to many, or many to one and one to many, in whichever way we like. When you are designing brains, it is important for you to know that in one of

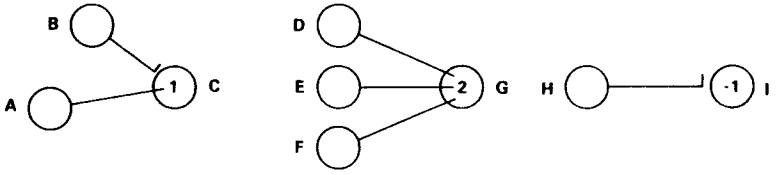

Figure 9

How threshold devices act on each other. Explanation of symbols: The circles stand for threshold devices. The L-shaped fiber between B and C stands for inhibition; the penetrating fiber from A to C means activation. Each active element contributes one unit of activation to the element (threshold device) to which it sends an activating connection. The threshold device becomes active when the activation reaches at least the threshold value indicated within the circle. An inhibitory connection from an active element subtracts 1 from the sum of all the units of activation reaching the same target element. A negative threshold (or threshold 0) implies activity in the absence of external activation. Such an element can be silenced by a corresponding amount of inhibition.

these threshold devices the output does not appear immediately upon activation of the input, but only after a short delay, say one tenth of a second. During this time the gadget performs its little calculation, which consists of comparing the quantity of its activation with its threshold.

You can already guess some of the things that a vehicle fitted with this sort of brain can do, but you will still be surprised when you see it in action. The vehicle may sit there for hours and then suddenly stir when it sights an olive green vehicle that buzzes at a certain frequency and never moves faster than 5cm/sec. Since our brand 5 vehicle is not interested in any other vehicles, you might say that the olive green vehicle is its special friend. You will have to conclude that Vehicle 5 has something like proper nouns in his mind, NAMES that refer to very particular objects, like James, Calcutta, or Jupiter.

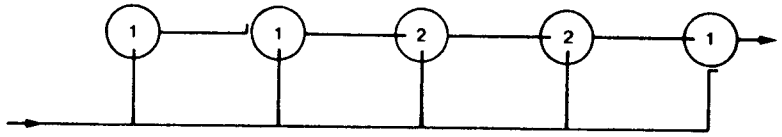

Figure 10a

A network that gives a signal when a burst of 3 pulses presents itself, preceded and followed by a pause.

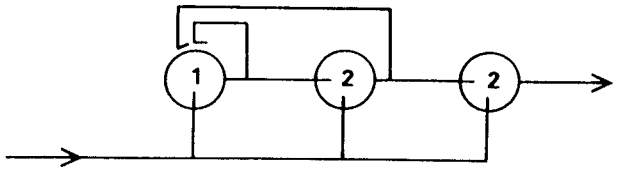

Figure 10b

A network of threshold devices that emits a pulse for every third pulse in a row in the input.

But Vehicle 5 can do much more than that. It can count (figure 10). It may associate only with groups of four vehicles, not more and not less, to make a party of five. Or it may visit every tenth source it encounters on its way. Or it may turn away from a vehicle whose number of sensors is a multiple of seven, implying that such vehicles bring bad luck. In some way, it seems to operate with NUMBERS.

If you fit such a vehicle with a very large number of shrewdly connected threshold devices, you may get it to play a passable game of chess. Or you may make it solve puzzles in LOGIC or prove theorems in euclidean geometry. You realize what I am driving at: with enough threshold devices it can do anything a computer can do, and computers can be made to do almost everything.

But where is the memory, some of you will ask, realizing that

most of the activities of a digital computer consist of putting data into memory, taking the data out again to perform some calculation, putting the results back into the memory, and so forth. The answer: there is room for memory in a network of threshold devices, if it is large enough. Imagine a threshold device connected to a sensor for red light. When it is activated by the red light, it activates another threshold device which in turn is connected back to the first device. Once a red light is sighted, the two devices will activate one another forever. Take a wire from the output of one of the two threshold devices and connect it to a bell: the ringing of the bell then signals the fact that at some time in the past this particular vehicle sailed in the vicinity of a source of red light.

This is an elementary sort of MEMORY. It is not difficult to understand how out of such elementary memory stores (consisting of reciprocally connected threshold devices) complex memories can be synthesized, with the possibility of storing extremely complex events. But there is a limit to the quantity of facts the vehicle can store this way. For instance, when storing numbers, if the vehicle has a bank of ten elementary memory devices, it cannot fit any number that has more than ten digits (in binary notation), since each elementary device can at most remember one digit by being active or inactive ("one bit of information").

There is a trick that can be used by our brand 5 vehicles to overcome the intrinsic limitation of their storage capacity. Imagine a vehicle involved in a calculation in which numbers occur that are much larger than the number of parts in the vehicle's own interior. You might think that such a task would be forever beyond the comprehension of that particular vehicle. Not so if we employ the following strategy. Let's transfer our vehicle to a large, sandy beach. The vehicle can crawl on the beach, leaving marks in the sand indicating the succession of digits in the large numbers that emerge from its calculations. Then it can crawl back, following

its own track, to read off the digits and put them back into the calculation.

The vehicle is never able to comprehend these large numbers at any one moment. But using itself as an instrument in a larger scheme involving the environment, and partly directed by it, it ends up with the correct result. (Of course, to be on the safe side, we must suppose that the sandy surface has no limits.) If you want a concrete example, think of the vehicle calculating the difference (small enough for it to comprehend) between two large numbers, which it can produce but not comprehend. It will produce one number by leaving marks on its way along the beach. It will produce the other number on its way back. And then it will measure the difference by counting the number of marks that are in default or in excess of the first number.

Later on, we will learn how to incorporate into a vehicle something quite analogous to the sand outside, and almost as boundless in its capacity.

Vehicle 6

Selection, the Impersonal Engineer

In this chapter things get slightly out of hand. You may regret this, but you will soon notice that it is a good idea to give chance a chance in the further creation of new brands of vehicles. This will make available a source of intelligence that is much more powerful than any engineering mind.

Out of the collection of vehicles that we have produced for the purposes of our experimentation, we will choose some of the more complicated specimens and put them onto a large table. Of course there will also be some sources of light, sound, smell, and so forth on the table, some of them fixed and some of them moving. And there will be various shapes or landmarks, including the cliff that signals the end of the table top.

Now you and I will gather a plentiful supply of materials (tin, plastic, threshold devices, wheels, motors, sensors, wires, screws and bolts) and proceed to build vehicles, taking as our models vehicles that we pick from the ones circulating on the table. Each time we copy a vehicle, we will put both the model and its copy back on the table, pick up another vehicle, copy it, and so on. Of

course we will not pick up vehicles that have fallen on the floor because they have proved their own inability to cope with the environment. We will be careful to produce vehicles at a pace that roughly matches the rate at which vehicles fall off the table, to prevent the race from dying out, on one hand, and to prevent the table from becoming unduly crowded, on the other.

Note that while we are playing this game, we won't have time to test the behavior or to study the wiring, let alone to understand the logic of the vehicles that we pick up as models for copying. Nor should we. All we are asked to do is to slavishly connect the parts according to the pattern in the model.

Note also that when we do this in a hurry, we are bound to make occasional mistakes. It may be our fault when our copy of a perfectly well-tested vehicle falls off the table as soon as we put it down. But it is also possible that we will unwittingly introduce a particularly shrewd variation into the pattern of connections, so that our copy will survive forever while the original may turn out to be unfit for survival after all.

It does seem surprising that errors arising in the sloppy execution of a task should act as germs for improvement. What is less astonishing is the creative power of a special sort of error consisting of new combinations of partial mechanisms, each of which is not disrupted in its own well-tested structure. This can easily happen when we pick up one vehicle as a model for one part of the brain and then by mistake pick up another vehicle as a model for another part of the brain. Such errors have a much greater chance of transcending the intelligence of the original plan.

This is an important point. If the lucky accidents live on forever, they will also have a multitude of descendants, for they will stay on the table all the time while the less lucky ones come and go. Therefore, they have a much greater chance of being picked up by the copyists as models for the next generation. Thus very good ideas

unwittingly introduced into the wiring, though improbable, do become quite widespread in the long run.

This story is quite old and goes by the name of Darwinian evolution. Many people don't like the idea that everything beautiful and marvelous in organic nature should be due to the simple cooperation of reproduction, errors, and selection. This is no problem for us. We have convinced ourselves that beautiful, marvelous, and shrewd machines can be made out of inorganic matter by this simple trick. Moreover, we already know that analysis is much more difficult than synthesis. Where there has been no conscious engineering at all, as in the case of our type 6 vehicles, analysis will necessarily produce the feeling of a *mysterious* supernatural hand guiding the creation. We can imagine that in most cases our analysis of brains in type 6 vehicles would fail altogether: the wiring that produces their behavior may be so complicated and involved that we will never be able to isolate a simple scheme. And yet it works.